

Classification of molecular sequence data using Bayesian phylogenetic mixture models

Elisa Loza-Reyes*, Merrilee Hurn[†] and Tony Robinson[‡]

July 28, 2011

Abstract

Rate variation among the sites of a molecular sequence is commonly found in applications of phylogenetic inference. Several approaches exist to account for this feature but they do not usually enable us to pinpoint the sites that evolve under one or another rate of evolution in a straightforward manner. In this paper we concentrate on phylogenetic mixture models as tools for site classification. Our method does not rely on prior knowledge of site membership to classes or even the number of classes. Furthermore, it does not require correlated sites to be next to one another in the sequence alignment, unlike some phylogenetic hidden Markov or change-point models. We present a simulation study to show that our approach is able to correctly classify the sites to evolutionary classes and we analyse the popular alignment of the mitochondrial DNA of primates. In both examples, all mixtures outperform commonly-used models of among-site rate variation and models that do not account for rate heterogeneity. Our method for site classification is directly relevant to the profiling of genes with unknown function, and its application may lead to the discovery of partitions not otherwise recognised in the alignment. In addition, we discuss computational aspects including

*Department of Biomathematics and Bioinformatics, Rothamsted Research, Harpenden, AL5 2JQ, UK.
Email: Elisa.Loza@rothamsted.ac.uk

[†]Department of Mathematical Sciences, University of Bath, Bath, BA2 7AY, UK. Email: M.A.Hurn@bath.ac.uk

[‡]Department of Mathematical Sciences, University of Bath, Bath, BA2 7AY, UK. Email: A.Robinson@bath.ac.uk

the use of simple Markov chain Monte Carlo (MCMC) moves to estimate phylogenetic models and argue that these move types can mix efficiently without tempering the MCMC chains.

Key words: among-site rate variation; allocation variable; Bayesian mixture model; Markov chain Monte Carlo; phylogeny

1 Introduction

Molecular phylogenetics is the inference and interpretation of evolutionary relations between taxa, typically different species or strains of viruses or bacteria, based on the taxa's DNA or protein sequences. The sequences are aligned on top of each other to form an alignment with as many rows as sequences observed, and roughly as many columns (or sites) as characters in the sequences. The conventional likelihood-based model for phylogeny inference (e.g. Felsenstein, 1981) contains three parameters of inferential interest: a tree graph that represents the evolutionary relations between the taxa; the branch lengths of this tree that measure the expected number of substitutions per site; and a stochastic process that models the evolution of the sequences along the branches of the tree (the latter is usually referred to as the evolutionary model). Such a model is complex but may still be too simple to capture important features of the generating process. In particular, it is not uncommon that sites under different functional constraints accumulate substitutions at different rates. It is now well understood that if rate variation among sites is present and not accounted for by the model, spurious parameter estimates can be produced (Huelsenbeck and Suchard, 2007 and references therein).

Various approaches have been proposed to account for among-site rate variation in phylogenetic inference, including the gamma model (Yang, 1993; 1994) and several more recent models involving finite mixtures of distributions (e.g. Pagel and Meade, 2004; Lartillot and Philippe, 2004; Huelsenbeck and Suchard, 2007; Webb, Hancock and Holmes, 2009). The latter type of models assume that a site is generated from a mixture of multiple processes, each of which may be indexed by a specific tree topology, a specific set of branch lengths and specific parameters of the stochastic evolutionary model.

Rate variation among sites may be related to quantitative differences in the rates of

substitution (e.g. sites with high rates versus sites with low rates) but also to qualitative differences in the pattern of substitution (e.g. sites with large transition/transversion rate ratios versus sites for which all substitution types occur at the same rate; Pagel and Meade, 2004). In phylogenetic applications it is possible to find quantitative among-site rate variation, qualitative variation, both or neither. Developments in phylogenetic mixture modelling have accounted for both types of rate variation and examples of this include Felsenstein and Churchill’s approach (1996). They account for quantitative variation in substitution rates among sites by a hidden Markov process that operates along the alignment assigning rates to sites from a finite pool of values. This method incorporates the biologically realistic assumption of correlation between the rates of evolution at consecutive sites, so that the chance of neighbouring sites evolving under the same rate is higher than that of distant sites. A disadvantage of this assumption, however, is that possible biases may be introduced by the removal of sites involving gaps in the alignment, or by other errors that result in consecutive observable sites not being direct neighbours in reality.

To model qualitative rate heterogeneity, Pagel and Meade (2004) use a Bayesian mixture of multiple stochastic evolutionary processes. Their model supposes that data at a given site arise from a mixture of multiple classes, each class indexed by a common-to-all-class tree and branch lengths, and a class-specific stochastic evolutionary process. The assumption of a common set of branch lengths across mixture components results in a phylogeny whose branches are a compromise over the possibly quite different substitution rates in the alignment. This may miss important substitutional heterogeneity and so Pagel and Meade (2008) and Meade and Pagel (2008) consider extensions to their original model, this time allowing for multiple sets of branch lengths.

Kolaczkowski and Thornton (2008) present a mixture similar to that of Meade and Pagel (2008), but conduct inference within a maximum-likelihood framework.

A related approach, called the CAT model (Lartillot and Philippe, 2004), considers qualitative mixtures of stochastic evolutionary processes which, for simplicity, all have the same set of substitution rates but different stationary probabilities. Inference on the number of classes is conducted using a Dirichlet process prior and the model is estimated via MCMC. In addition to moving in the space of usual phylogenetic parameters (e.g. tree topology, branch lengths), the MCMC sampler developed by Lartillot and Philippe also moves in the space of

number of classes, jumping between mixtures of different dimensions as the run proceeds.

Huelsenbeck and Suchard (2007) consider quantitative mixtures of branch lengths in which sites are partitioned into classes according to a Dirichlet process prior. Sites that are assigned to the same class share a common set of branch lengths, while all sites, irrespective of their class, share a common tree and stochastic evolutionary model. Both the number of classes and the assignment of sites to classes are treated as random variables and, together with the usual phylogenetic parameters, are objects of inferential interest.

One aspect of mixture models that has been under-explored in the phylogenetics literature is their use for site classification through the introduction of latent allocation variables. The allocation variables identify the underlying class of a site and thus enable us to decompose the complicated structure of a mixture into simpler structures. In a phylogenetic context, mixture components may have a direct biological interpretation and site classification can lead to insights of structure and heterogeneity in the alignment that are not otherwise easily uncovered.

The purpose of this study was, therefore, to extend the functionality of phylogenetic mixture models to include allocation variables and investigate their use for site classification. Lartillot and Philippe (2004) and Huelsenbeck and Suchard (2007) incorporate allocation variables, but straightforward statements about site classification are obscured by the estimation of the number of mixture components. Their MCMC samplers move in a space of mixtures of different dimensions, and so sites get allocated to an ever-varying number of components as the MCMC run progresses. The mixture in Pagel and Meade (2004) does not incorporate allocation variables and site classification therefore involves *a posteriori* processing of their analysis output.

In our study we do not attempt inference on the number of mixture components but, instead, consider mixtures with fixed dimensionality. We develop a heuristic technique, based on Bayes factors and inspection of the analysis output, to select the number of mixture components and the type of mixture that best fits the data. We demonstrate model selection and site classification through applications to both synthetic and real DNA data. The results suggest that our method is able to detect heterogeneity in these data and classify the sites to evolutionary components with high accuracy.

With regard to model estimation, we consider commonly-used MCMC move types that

update tree topology and branch lengths *en bloc* and argue that this may have detrimental effects on the estimation performance of the sampler. We present an alternative set of move types to update the parameters of a mixture or non-mixture phylogenetic model, and investigate their performance. We show that our algorithm achieves the same, or greater, efficiency than existing methods with potential for a reduction in computational cost.

2 Bayesian phylogenetic mixtures for site classification

2.1 The models

The backbone of likelihood-based phylogenetic methods is a *homogeneous model* that posits that the characters at a site in a DNA alignment are an independent realisation of a continuous-time Markov process, with state space $\mathcal{I} = \{A, C, G, T\}$, that evolves on the branches of a bifurcating tree topology, ϕ , and has realisations at the leaves of this tree. The instantaneous rate matrix, Q , that generates the Markov process is indexed by a (possibly vector) parameter θ . There are several proposed parametrisations of the Q -matrix in the literature (e.g. Jukes and Cantor, 1969; Hasegawa et al., 1985) with the most general time-reversible one called the GTR matrix (Lanave et al., 1984; Tavaré, 1986), where

$$Q(\theta) = \begin{pmatrix} q_{AA} & r_{AC}\pi_C & r_{AG}\pi_G & r_{AT}\pi_T \\ r_{AC}\pi_A & q_{CC} & r_{CG}\pi_G & r_{CT}\pi_T \\ r_{AG}\pi_A & r_{CG}\pi_C & q_{GG} & r_{GT}\pi_T \\ r_{AT}\pi_A & r_{CT}\pi_C & r_{GT}\pi_G & q_{TT} \end{pmatrix} \quad (1)$$

and $\theta = (r, \pi)$ is a collection of six substitution rates $r = (r_{AC}, \dots, r_{GT})$ and four stationary probabilities $\pi = (\pi_A, \dots, \pi_T)$ with constraints $r_m, \pi_i \geq 0$ and $\sum r_m = \sum \pi_i = 1$ ($m = AC, \dots, GT; i = A, \dots, T$). The off-diagonal values of Q are non-negative while the diagonal entries are defined so that each row adds up to zero. This matrix is usually standardised so that $\sum_{i \in \mathcal{I}} -q_{ii}\pi_i = 1$ (Felsenstein, 1981) and, as a result, the branch lengths are a measure of the expected number of substitutions per site in the molecular sequence alignment. The Markov process of character substitution is time-reversible, a feature that prevents us from inferring rooted trees. Thus, for an observed alignment of size S sequences $\times N$ sites, parameter ϕ takes values in the set of unrooted bifurcating leaf-labelled trees for S taxa;

branch lengths are real valued; and the space in which parameter θ takes values is dictated by the chosen parametrisation of matrix Q . The objective of the analysis is usually inference about the tree topology, ϕ , this tree's branch lengths (denoted by a set $t = \{t_1, \dots, t_{2S-3}\}$) and θ .

Building upon the homogeneous model, we account for among-site rate variation using a finite mixture of distributions of the type

$$x_n \mid \omega, \phi, t, \theta, k \sim \sum_{j=1}^k \omega_j p(x_n \mid \phi, t_j, \theta_j), \text{ independently for } n = 1, \dots, N, \quad (2)$$

where x_n is the character vector at site n ; k is the number of mixture components; $\omega = (\omega_1, \dots, \omega_k)$ are the mixture proportions ($\omega_j \geq 0$ and $\sum_{j=1}^k \omega_j = 1$); each component j ($j = 1, \dots, k$) has set of branch lengths t_j and parameters of the Q -matrix θ_j collectively denoted by $t = (t_1, \dots, t_k)$ and $\theta = (\theta_1, \dots, \theta_k)$; and $p(x_n \mid \phi, t_1, \theta_1), \dots, p(x_n \mid \phi, t_k, \theta_k)$ are the k component likelihoods. Model (2) thus asserts that characters at site n are generated from a mixture of k different evolutionary components occurring in proportions $\omega_1, \dots, \omega_k$. To decompose the structure of this mixture, a set of latent allocation variables, $z = (z_1, \dots, z_N)$, is introduced where each $z_n \in \{1, \dots, k\}$ is such that

$$x_n \mid \phi, t, \theta, k, z_n \sim p(x_n \mid \phi, t_{z_n}, \theta_{z_n}), \text{ independently for } n = 1, \dots, N. \quad (3)$$

This formulation not only accounts for both quantitative and qualitative rate heterogeneity, but also provides a means to class discovery by the use of z . In addition to classifying the sites to evolutionary components, mixture (2) also enables us to discern the profiles of each class by estimating the component-specific parameters. So, the analysis may lead to statements such as ‘class 1 is more conserved than class 2 as the former displays a shorter total branch length than the latter’ or ‘the nucleotide composition of the two classes is quite different, as reflected by the estimated stationary probabilities’.

The joint prior of all parameters can be expressed as

$$\begin{aligned} p(\omega, z, \phi, \theta, t) &= p(\omega)p(z \mid \omega)p(\phi \mid z, \omega)p(\theta \mid \phi, z, \omega)p(t \mid \theta, \phi, z, \omega) \\ &= p(\omega)p(z \mid \omega)p(\phi)p(t)p(\theta) \end{aligned} \quad (4)$$

where we have suppressed the explicit conditioning on k because we consider only mixtures with a fixed number of components and make independence assumptions between all parameters other than z and ω . In our study, the prior for ω is taken to be the symmetric Dirichlet distribution $\omega \sim \text{Dir}_k(\rho, \dots, \rho)$ (and we generally set $\rho = 3$). Conditional on ω , the allocations z_1, \dots, z_N are assumed independent and identically distributed

$$\Pr(z_n = j \mid \omega) = \omega_j, \quad j = 1, \dots, k. \quad (5)$$

We make the following standard choices for the priors on phylogenetic parameters. All tree topologies are assumed to be equally likely *a priori*; that is, we take a discrete uniform prior for ϕ . The prior distribution for branch lengths makes an assumption that the $2S - 3$ branches for each of the k components behave independently both within components and across components. Exponential priors on individual branch lengths are specified, with exponential-rate parameter β (and we set $\beta = 10$ in line with similar published models) so that $\mathbf{E}(t_{h,j}) = 1/\beta$ for branch length h in the j th mixture component ($h = 1, \dots, 2S - 3$; $j = 1, \dots, k$). For the parameter vectors θ of the k instantaneous rate matrices, we assume independent prior distributions on each r_j and π_j of the form $r_j \sim \text{Dir}_6(1, \dots, 1)$ and $\pi_j \sim \text{Dir}_4(1, \dots, 1)$.

Throughout, the model specified by (2) and (3) is referred to as the $Q + t$ mixture model.

In our examples, we will also consider nested submodels of the $Q + t$ mixture. Firstly, we consider mixtures of multiple Q matrices which share a common set of branch lengths, t , and tree topology, ϕ , (Pagel and Meade, 2004):

$$x_n \mid \omega, \phi, t, \theta, k \sim \sum_{j=1}^k \omega_j p(x_n \mid \phi, t, \theta_j), \quad \text{independently for } n = 1, \dots, N. \quad (6)$$

Restricting this further we also consider mixtures of branch lengths and Q matrices, but where the Q matrices across components share the same stationary probabilities, i.e. $\theta_1 = (r_1, \pi), \dots, \theta_k = (r_k, \pi)$:

$$x_n \mid \omega, \phi, t, \theta, k \sim \sum_{j=1}^k \omega_j p(x_n \mid \phi, t_j, r_j, \pi), \quad \text{independently for } n = 1, \dots, N. \quad (7)$$

Both models can be augmented with allocation variables. We refer to model (6) and its corresponding augmented formulation as the Q mixture, and to model (7) and its augmented

version as the $r + t$ mixture.

2.2 Likelihood computation

The likelihood function under the most general $Q + t$ mixture is the product of the distributions at individual sites (equation (3)), from site 1 to N :

$$L(\phi, t, \theta | x, z) = \prod_{n=1}^N p(x_n | \phi, t_{z_n}, \theta_{z_n}). \quad (8)$$

We assume that substitutions at different branches of the tree and among different sites in the alignment are independent of one another. Likelihood (8) is usually computed for a specific tree and so each tree topology requires a reformulation of this function according to its corresponding branching structure; the larger the tree the more computationally prohibitive the calculation. A recursive technique for the efficient computation of phylogenetic likelihood functions, called the pruning algorithm, was introduced by Felsenstein (1981), and this is the algorithm that we use.

2.3 Model fitting

Markov chain Monte Carlo (MCMC) will be required to fit models of this complexity and we present the basic move types in our MCMC sampler. A distinctive feature of our method is that changes to the topology are separated from those in branch lengths. This is particularly important for some of the mixtures, where the components share a common topology but have different sets of branch lengths. The set of move types that we use is:

- (a) updating the tree topology ϕ ;
- (b) updating branch lengths t ;
- (c) updating the substitution rates r ;
- (d) updating the stationary probabilities π ;
- (e) updating the mixture proportions ω ;
- (f) updating allocation z_n .

One complete pass over these six moves is an iteration, the basic time step of our algorithm. The first two move types focus on the tree while the next two concentrate on the parameters

of the models on the tree; the last two move types concern the mixture allocations and proportions. We now consider the three groups separately in the context of the most general $Q + t$ mixture model, before investigating their performance later in the paper.

2.3.1 Updating the tree topology and branch lengths

The tree topology is updated via the nearest neighbour interchange (NNI) (Robinson, 1971; Moore, Goodman and Barnabas, 1973), in which one of the two nearest neighbours of the current topology (in NNI space) is proposed with equal probability. NNI generates a candidate topology while preserving the current set of branch lengths. A separate proposal mechanism is then used to update branch lengths while maintaining the same topology. We consider two different proposals for branch lengths:

- Branch length multiplier (BLM). Also known as proportional shrinking and expanding (Yang, 2006), this proposal updates the length of a randomly chosen branch $t_{h,j}$ by multiplying it by a quantity generated from the density

$$f(m) = (\lambda m)^{-1}, \quad 1/\delta < m < \delta \quad (9)$$

where $\lambda = 2 \log \delta$ and $\delta > 1$ acts as a tuning parameter.

- Branch length normal additive (BLNA). Also known as the sliding window proposal (Huelsenbeck and Ronquist, 2001), this mechanism updates a randomly chosen branch length $t_{h,j}$ via an additive Gaussian perturbation, $t'_{h,j} \sim N(t_{h,j}, \sigma^2)$, so that σ^2 acts as the tuning parameter. If negative branch lengths are proposed, they are reflected at zero with the proposal still remaining symmetric.

BLM may be thought of as self-tuning as the variance of the proposed branch length is proportional to the square of the original length. This works well when exploring large branches but can be a bit sticky when branch lengths are small as it can take a large number of iterations to move a short distance. On the other hand, a candidate branch length generated from the BLNA proposal has a step size which depends only on the tuning parameter σ^2 and not on the current branch length. This makes it hard for BLNA to work equally effectively at both large and small scales. In experiments, we achieved best performance by alternating between BLM and a BLNA tuned for small branch lengths (Supplementary Material I).

The acceptance probability of a branch length proposed from either BLM or BLNA is

$$\alpha(t_{h,j}, t'_{h,j}) = \min \left\{ 1, \frac{p(t'_{h,j})}{p(t_{h,j})} \frac{L(\phi, t', \theta | x, z)}{L(\phi, t, \theta | x, z)} \frac{q(t'_{h,j}, t_{h,j})}{q(t_{h,j}, t'_{h,j})} \right\} \quad (10)$$

which simplifies to $m e^{-\beta(t'_{h,j} - t_{h,j})}$ for the BLM proposal and to $e^{-\beta(t'_{h,j} - t_{h,j})}$ for BLNA, times the likelihood ratio in both cases.

2.3.2 Updating the Markov process parameters

The j th component of the $Q + t$ mixture has a set of parameters controlling the substitution rates plus a set of stationary probabilities, $r_{AC,j}, \dots, r_{GT,j}$ and $\pi_{A,j}, \dots, \pi_{T,j}$, respectively. Since we can treat each mixture component separately for updating purposes, we drop the subscript j . Both types of parameters are constrained to sum to one and, as they utilise the same type of proposal, here we concentrate on the substitution rates.

We suggest generating a new set of substitution rates, r' , from a Dirichlet distribution centred at the current rate values with a positive shift $\epsilon > 0$ and with tuning parameter $\alpha > 0$; i.e. $r' \sim \text{Dir}_6(\alpha(r_{AC} + \epsilon), \dots, \alpha(r_{GT} + \epsilon))$. The variance of the m th element of a rate vector proposed with this move type, henceforth referred to as the ϵ Dirichlet proposal, is:

$$\text{var}(r'_m) = \frac{\alpha(r_m + \epsilon)(\alpha_0 - \alpha(r_m + \epsilon))}{\alpha_0^2(\alpha_0 + 1)} \quad (11)$$

where $\alpha_0 = \sum_{m=AC}^{GT} \alpha(r_m + \epsilon) = \alpha(1 + 6\epsilon)$. If $\epsilon = 0$ (Larget and Simon, 1999), when r_m is close to zero so too is $\text{var}(r'_m)$. This can create a cycle in which the MCMC sampler keeps proposing candidate rates very close to zero because the step size of the proposal is nearly zero, typically needing many iterations to escape. We further investigate this behaviour in a later section.

2.3.3 Updating the mixture parameters

Updating the allocation variables and the vector of mixture proportions is a fairly standard problem in mixture modelling. The usual approach is to update the allocations one at a time using either a uniform Metropolis or a Gibbs proposal (the number of mixture components k tends to be small in our applications so this is feasible). The mixture proportions are usually updated using a Gibbs sampler since their posterior conditional is easily seen to be

a Dirichlet distribution with parameters $\rho + N_1, \dots, \rho + N_k$, where $N_j = \sum_{n=1}^N I[z_n = j]$ is the number of sites allocated to component j and $I[\cdot]$ is the indicator function. A well known difficulty of this combination of proposals is that they may mix badly when one or more components become quite small or when the other parameters characterising the components make it hard for a site to swap components (see Leslie, 2007 or Hurn et al., 2008 for examples in quite different application areas). In the latter case, Leslie (2007) and Hurn et al. (2008) both suggest a joint updating strategy. However here we are primarily worried about the former case. Given our experience in updating r and π , we again propose using a shifted Dirichlet approach, here replacing the Gibbs draw from a $Dir_k(\rho + N_1, \dots, \rho + N_k)$ by a Metropolis-Hastings proposal, $\omega' \sim Dir_k(\rho + N_1 + \epsilon, \dots, \rho + N_k + \epsilon)$ with $\epsilon > 0$. The acceptance probability of this move type simplifies to

$$\alpha(\omega, \omega') = \min \left\{ 1, \prod_{j=1}^k (\omega_j / \omega'_j)^\epsilon \right\} \quad (12)$$

and so a high acceptance rate is maintained for small values of ϵ .

2.4 Model selection

Turning to the decision of choosing which model to use for a particular set of data, Bayes factors can be used to summarise the evidence provided by the data in favour of one model relative to another (Kass and Raftery, 1995). When each model is equally likely *a priori*, the Bayes factor is defined as the ratio of the marginal likelihood under model M_1 to the marginal likelihood under a second model, M_0 , given the data, x . The marginal likelihood for model M_i is the expectation under the prior of the likelihood of the data x , all conditioned on the model M_i (or, equivalently, the integral over the parameters of the joint distribution of the data and the prior conditioned on the model),

$$p(x|M_i) = \int_{\vartheta_i} p(x|\vartheta_i, M_i) p(\vartheta_i|M_i) d\vartheta_i \quad (13)$$

where ϑ_i is the parameter vector of model M_i . Bayes factors are usually interpreted on the log scale using the rule of thumb that $2 \ln(BF) > 10$ indicates very strong evidence in favour of one of the models, $0 \leq 2 \ln(BF) \leq 2$ indicates no significant difference between the models, and with a range of levels in between according to a scale provided in Kass

and Raftery (1995). In examining more than two models, the values of the log marginal likelihoods may be compared directly rather than forming multiple Bayes factors.

The integral in Equation (13) is typically intractable, except for the most elementary phylogenetic applications, and its estimation is the topic of considerable interest. A commonly used estimator is the harmonic mean estimator of Newton and Raftery (1994). This is a form of importance sampling estimator, taking the posterior as its importance distribution which means it can be calculated from the MCMC chain used for fitting the model at little extra cost. The harmonic mean estimator is known to be relatively unstable although studies suggest that it is accurate enough for interpretation on the log scale (Kass and Raftery, 1995 and references therein). A recent paper by Xie et al., (2011) details its shortcomings, strongly suggesting instead a much more computationally expensive estimation approach. Given the already computationally intensive nature of our model fitting, we have chosen to work with the harmonic mean estimator since we are interested in log marginal likelihoods. However we do bear in mind its potential instability when choosing run lengths.

3 Classification of simulated data

To demonstrate the key features of our approach we generated a synthetic DNA alignment of size 16 sequences \times 2 500 sites, with the software package Seq-Gen (Rambaut and Grassly, 1997). Sites 1 – 1500 were generated from an evolutionary class with substitution rates $r_1 = \{r_{AC} = 0.0500, r_{AG} = 0.4000, r_{AT} = 0.0500, r_{CG} = 0.0500, r_{CT} = 0.4000, r_{GT} = 0.0500\}$, stationary probabilities $\pi_1 = \{\pi_A = 0.3220, \pi_C = 0.3040, \pi_G = 0.1080, \pi_T = 0.2660\}$ and total branch length $T_1 = 10$. Sites 1501 – 2500 were simulated with $r_2 = \{r_{AC} = 0.1009, r_{AG} = 0.3645, r_{AT} = 0.1506, r_{CG} = 0.0639, r_{CT} = 0.3044, r_{GT} = 0.0157\}$; $\pi_2 = \{\pi_A = \dots = \pi_T = 0.2500\}$ and $T_2 = 0.1$. We arbitrarily labelled the former as class 1 and the latter as class 2. Both classes were generated under the same tree topology and this was randomly sampled from the space of all unrooted bifurcating trees that relate 16 sequences. In our experiments, the topology was held fixed at the generating topology.

The intention here is to assess whether the classification method is able to detect the substitutional differences between the two classes and correctly allocate sites to evolutionary groups without prior knowledge of the partitioning.

Before the runs for inference, we conducted several exploratory runs to tune the MCMC proposals. We estimated each model at least twice and verified that each of these independent runs converged to the same region of the parameter space. Unless otherwise stated, the runs for inference comprised 60 000 iterations of our MCMC sampler, and we discarded the first quarter as burn-in.

Figure 1 shows the estimated log marginal likelihoods for a range of mixture models; $r+t$, Q and $Q+t$ with varying number of components. The log-likelihood for $k=1$ is common across all mixture types and it corresponds to a fit of the data with the homogeneous model; it is clear that any mixture performs better than the homogeneous model. For comparison, we also fitted a model in which the rates of substitution are allowed to follow a gamma distribution with four discrete categories (Yang, 1993; 1994). The log marginal likelihood of the discrete-gamma model, as this method is usually known, was estimated in $-29\,581$. All two-component mixtures perform better than the discrete-gamma model which suggests that the substitutional heterogeneity in the data can only be adequately explained by multiple sets of branch lengths and more than one set of rate parameters. As expected, the log marginal likelihood reaches a plateau with a two-component $Q+t$ mixture and continues to show improvements with some three- and four-component mixtures. A typical output from Q and $Q+t$ mixtures with $k>2$ displays partitions of the sites that are based on nucleotide composition (e.g. a component rich in one nucleotide type with the corresponding complementary component that is poor in that nucleotide), and which may be driving the rise in log-likelihood. Although the partition of sites by character composition is reasonable, we are interested in discovering class structure beyond the obvious differences. But perhaps more importantly, substitution rates cannot be reliably estimated for components that are poor in one or more nucleotide types. A fit with a $r+t$ mixture, on the other hand, shows a negligible increase in log likelihood as the number of components augments. Indeed, $r+t$ mixtures of sizes three and four show nearly empty components with very low weights. This suggests that, whenever the partition of the sites is not based on character composition, as in an $r+t$ fit, two components are sufficient to describe the data. At two components, there is very strong evidence for a $Q+t$ mixture ($2\ln(BF_{Q+t\,vs\,Q}) = 1\,462$ and $2\ln(BF_{Q+t\,vs\,r+t}) = 1\,196$), and this is the model we use to discuss site classification.

MCMC proposals for the classification run were tuned to values $\delta = 1.5$ for the BLM




Figure 1: Estimated log marginal likelihoods for the models fitted to the synthetic DNA alignment. MCMC runs for mixtures with four components comprised 20 000 iterations, with the first quarter discarded as burn-in. Plotted data: log-likelihood for a homogeneous model: $-32\,219$; $Q + t$ mixture with 2–4 components: $-27\,368$, $-27\,134$, $-26\,467$; Q mixture with 2–4 components: $-28\,099$, $-27\,563$, $-26\,530$; $r + t$ mixture with 2–4 components: $-27\,966$, $-27\,908$, $-27\,869$; discrete-gamma model with 4 categories: $-29\,581$.

move; $\sigma = 0.08$ for BLNA; $\alpha_r = 900$ and $\alpha_\pi = 700$ for the ϵ Dirichlet proposal for substitution rates and stationary probabilities, respectively; $\epsilon_\theta = 0.0001$ to correct the ϵ Dirichlet move and $\epsilon_\omega = 0.0001$ to correct the proposal for mixture proportions.

Figure 2 shows the estimated posterior probabilities of classification to the first component (A details how these probabilities were obtained). High probabilities coincide with sites originally generated from class 1 while low probabilities match the positions generated from the second class. The crossover at which sites were simulated from different evolutionary classes is strikingly well recovered by the method. Moreover, ergodic averages for the remaining parameters in the mixture coincide favourably with the generating values.

4 Classification of mitochondrial DNA

In a second application, we analysed an alignment of mitochondrial DNA (mtDNA) sequences from the primate species human; gorilla; chimpanzee; orangutan; gibbon; crab-eating macaque; common squirrel monkey; Philippine tarsier and ring-tailed lemur (Brown




Figure 2: Estimated posterior probabilities of classification for the synthetic DNA alignment. High and low probabilities clearly demarcate the two classes used in generating the data.

et al., 1982; Hayasaka et al., 1988). This alignment, of size 9 sequences \times 888 sites, comprises the portions of two protein-coding genes (sites 1 – 694) and a transfer RNA (tRNA) region (sites 695 – 888). Transfer RNA is a highly conserved molecule in charge of translating the information encoded by coding genes into the protein alphabet. Such a translation process is achieved by mapping each set of three consecutive, non-overlapping DNA characters within a coding region into one amino acid. A coding DNA triplet is called a codon, and the second position (*cp2*) of a codon is known to undergo substitutions at slower rates than the first (*cp1*) and third codon positions (*cp3*; Fitch and Markowitz, 1970). This difference in substitution rates relates to the fact that a change at the third codon position does not always affect the resulting protein but a change at *cp2* may, more likely, alter the final product and result in a mutation.

In this analysis, we are interested in detecting the evolutionary heterogeneity that exists between the different codon positions and the tRNA region. We also want to estimate the component-specific parameters in order to understand the features of the model that best discriminate between classes.

The primate mtDNA alignment has been analysed extensively using phylogenetic methods (e.g. Yang, 1995; Larget and Simon, 1999; Suchard et al., 2001), in most cases assuming four evolutionary classes (corresponding to the three codon positions plus the tRNA region). Most of these previous approaches have relied on prior knowledge about site membership which may be restrictive and error prone. For instance, in a study by Yang (1995), some

sites within the tRNA region were *a priori* misclassified resulting in inaccurate parameter estimates, as stated in the mtprim9.nuc file of the software package PAML4 (Yang, 2007).

We fitted $Q + t$ and Q mixtures with different number of components to the primate mtDNA alignment. Figure 3 shows the log marginal likelihoods estimated for each model from 45 000 iterations, following a burn-in period of 15 000 iterations. It is clear that the data contain heterogeneity that is not fully accounted for by either the homogeneous or the discrete-gamma models. A $Q + t$ mixture with two components improves upon the homogeneous model by 576 log-units and upon the discrete-gamma model by 414 log-units. A Q mixture with two components improves upon the homogeneous and discrete-gamma models even more, and greater improvements are observed as the number of mixture components augments for these two types of models. At first sight, the results point to the Q mixture with four components as the best model. However, a closer look at the analysis output reveals that the partition of sites into mixture components is, in some cases, mainly driven by the nucleotide composition of the sites. Figure 4 shows the ergodic averages of stationary probabilities obtained from the analyses of the data with Q and $Q + t$ mixtures with one to four components. All mixtures other than the $Q + t$ with two components, contain a component that is poor in A nucleotides.

To investigate further the partition of sites by nucleotide composition, we fitted the primate mtDNA alignment with a $r + t$ mixture. The log marginal likelihoods for a mixture of this type are shown in Figure 3. When the mixture components are not indexed by component-specific π s (in contrast with both the Q and $Q + t$ mixtures), the classification process is unable to separate the sites by their character content and the increase in log likelihood is more gradual. In fact, $r + t$ mixtures of sizes three and four show poorly estimated components with very low weights which suggests that two components are enough to account for the rate variation in this alignment. Collectively, these results indicate that the large increase in log likelihoods for $Q + t$ and Q mixtures of sizes three and four, relative to those with two components, is mainly due to uninteresting partitions of the sites by nucleotide composition. Figure 4 shows that this is also the case for an analysis with a two-component Q mixture. We therefore decide on the $Q + t$ mixture with two components to discuss site classification.

MCMC proposals for the classification run were tuned from several exploratory runs to




Figure 3: Estimated log marginal likelihoods for the models fitted to the primate mtDNA alignment. The MCMC chain for mixtures with four components comprised 15 000 samples after discarding the initial 5 000 as burn-in. Plotted data: log-likelihood for a homogeneous model: $-5\,201$; $Q + t$ mixture with 2–4 components: $-4\,625$, $-4\,412$, $-4\,300$; Q mixture with 2–4 components: $-4\,564$, $-4\,404$, $-4\,243$; $r + t$ mixture with 2–4 components: $-4\,710$, $-4\,620$, $-4\,557$; discrete-gamma model with 4 categories: $-5\,039$.

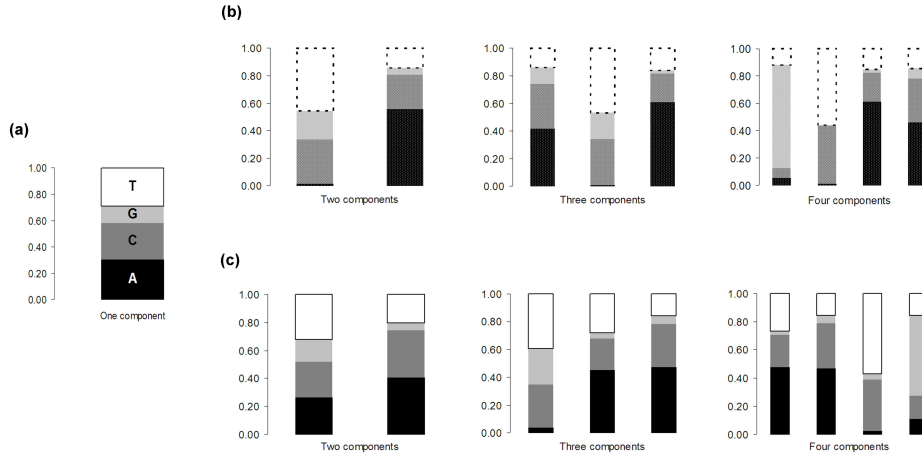


Figure 4: Ergodic averages of stationary probabilities from a fit of the primate mtDNA alignment with (a) a homogeneous model; (b) a Q mixture with two to four components; and (c) a $Q + t$ mixture with two to four components.

values $\delta = 1.5$ for the BLM move; $\sigma = 0.06$ for BLNA; $\alpha_r = 800$ and $\alpha_\pi = 600$ for the ϵ Dirichlet proposal for substitution rates and stationary probabilities, respectively; $\epsilon_\theta = 0.0001$ to correct the ϵ Dirichlet move and $\epsilon_\omega = 0.0001$ to correct the proposal for mixture proportions.

Table 1 reports the ergodic averages of model parameters and the estimated integrated autocorrelation times, $\hat{\tau}$ (B gives more details about τ). The ergodic averages for the rates of substitution agree with the bias that favours transitions (a substitution from $A \rightarrow G$ or $C \rightarrow T$) over transversions (any other substitution). The two components are well-differentiated by their branch lengths; the second class evolves under a tree with a total branch length around ten times longer than that of the first component. Branch lengths are the quantities that best discriminate between components and so the inclusion of multiple sets of branch lengths seems to be reasonable.

Figure 5 shows the estimated posterior classification probabilities of sites belonging to the second component (labelled as $j = 2$ in Table 1). For ease of visual interpretation, we have rearranged the protein-coding genes according to codon positions; sites 1 – 232 correspond to *cp1*, sites 233 – 463 to *cp2* and sites 464 – 694 to *cp3*, but there is nothing in the formulation of the classification method that requires such a rearrangement. Sites with high posterior probabilities in Figure 5 mostly occur within the *cp3* and *cp1* regions and, previously, we found that the second component has substantially larger branch lengths than component 1. This profile agrees with the theory; the *cp1* and *cp3* regions accumulate substitutions at higher rates than the more conserved *cp2* and tRNA, which is captured by both the classification probabilities and the parameter estimates.

The consensus tree topology, obtained as the 50% majority-rule, is shown in Figure 6 and it matches the published topologies in Yang (1995), Larget and Simon (1999) and Suchard et al. (2001).

5 Performance of the MCMC moves

We first compare our separated topology/branch length moves with the combined topology and branch length LOCAL proposal (Larget and Simon, 1999), which is implemented in the phylogenetic software package MrBayes (Huelsenbeck and Ronquist, 2001) and commonly




Figure 5: Estimated posterior classification probabilities for the primate mtDNA alignment, from an analysis with a two-component $Q+t$ mixture. High classification probabilities occur at the *cp1* and *cp3* regions, which are known to undergo substitutions at higher rates than the highly-conserved *cp2* and tRNA segments. The effect of horizontal lines is due to the presence of sites with exactly the same character composition.




Figure 6: The 50% majority-rule consensus topology obtained from the chain of sampled topologies during the analysis of the primate mtDNA alignment with a two-component $Q+t$ mixture. The posterior support for clades is shown on the branches.

<i>parameter</i>	<i>j=1</i>		<i>j=2</i>	
	<i>ergodic average</i>	$\hat{\tau}$	<i>ergodic average</i>	$\hat{\tau}$
$r_{(AC,j)}$	0.0332	205	0.0363	46
$r_{(AG,j)}$	0.3658	118	0.5362	190
$r_{(AT,j)}$	0.0622	58	0.0334	64
$r_{(CG,j)}$	0.0413	79	0.0617	149
$r_{(CT,j)}$	0.4791	95	0.2922	177
$r_{(GT,j)}$	0.0182	64	0.0399	200
$\pi_{(A,j)}$	0.2633	69	0.4044	134
$\pi_{(C,j)}$	0.2532	40	0.3390	88
$\pi_{(G,j)}$	0.1573	21	0.0537	73
$\pi_{(T,j)}$	0.3262	38	0.2029	131
ω_j	0.5534	31	0.4466	31
$\sum int_j$	0.1056	34	1.0941	28
$\sum ext_j$	0.2725	58	3.0993	43

Table 1: Ergodic averages of model parameters and estimated integrated autocorrelation times, $\hat{\tau}$, from an analysis of the primate mtDNA alignment with a two-component $Q + t$ mixture. Here j indexes the mixture components and the notation $\sum int_j$ and $\sum ext_j$ refers to the total length of interior and exterior branches for the j th component.

used for phylogenetic inference (e.g. Lartillot and Philippe, 2004; Nylander et al., 2004). Poor performance of LOCAL has been documented before for applications where only a few trees are supported by the data (e.g. Lakner et al., 2008) and our simulations agree with this observation. We consider a synthetic DNA alignment of 6 sequences \times 2500 sites generated using the program Seq-Gen (Rambaut and Grassly, 1997) with parameters: $r = \{r_{AC} = 0.140, r_{AG} = 0.340, r_{AT} = 0.090, r_{CG} = 0.008, r_{CT} = 0.420, r_{GT} = 0.002\}$; $\pi = \{\pi_A = \dots = \pi_T = 0.25\}$ and $(\phi, t) = (((Taxon2 : 0.16, Taxon4 : 0.34) : 0.61, Taxon5 : 0.2) : 0.53, (Taxon3 : 0.48, Taxon6 : 2.14) : 0.35, Taxon1 : 0.05)$. A homogeneous model was fitted with the substitution rates and stationary probabilities held fixed at values very close to their generating values. We extended the comparison to include MrBayes' inbuilt

implementation of Metropolis-coupled MCMC ((MC)³; Geyer, 1991) which is a technique to improve the movement of an MCMC sampler by coupling together a series of successively easier-to-sample distributions, but at a correspondingly higher computational cost. We opted for MrBayes’ default tuning of the LOCAL proposal and used two tempered chains when applying (MC)³.

Figure 7 shows the trace plots and corresponding autocorrelation functions monitoring one aspect of the output, the total branch length T (we use the total rather than individual lengths as interior branches are not consistently labelled across the different topologies visited during an MCMC run). It is clear that Metropolis-coupling the LOCAL proposal improves performance (compare Figure 7(b) with (d)). It is even clearer, however, that the separate topology/branch length moves are the best choice here (see Figure 7(f)). We can quantify the comparison of the methods by estimating the integrated autocorrelation time, τ (e.g. Green and Han, 1992); τ gives the factor by which a particular MCMC algorithm increases the variance of ergodic averages over those obtained by independent sampling. The smaller τ , the better the MCMC algorithm. In this example, the MCMC LOCAL algorithm gives $\hat{\tau} = 115.54$, (MC)³ LOCAL has $\hat{\tau} = 69.21$ while our proposed combination of moves has $\hat{\tau} = 12.01$. However, any reduction in $\hat{\tau}$ should be weighed against the computational costs of the different approaches. The MCMC LOCAL took 110 seconds, (MC)³ LOCAL 201 seconds while our proposed combination of moves took considerably more at 1740 seconds. In other words, we can decrease the variance by a factor of $115.54/12.01 = 9.62$ by using our proposals instead of LOCAL but it will take $1740/110 = 15.8$ times longer. This suggests it would be better simply to run the less efficient LOCAL for more iterations. However this is a hard comparison to make entirely fairly because MrBayes is a commercial package while our code is not; a speed-up of 2 in our code would reverse this conclusion. We also note here that the Metropolis-coupling of LOCAL decreases the variance by a factor of just 1.67 while increasing the time taken by a factor of 1.83.

One reason for the poor performance of LOCAL is that this proposal generates a new set of branch lengths and, as a by-product, it may also generate a new topology as part of the same updating step. The LOCAL proposal can make large moves in topology space, but when this move is not supported by the posterior distribution both the candidate topology and any change in branch lengths will be rejected. Therefore, when the data support only

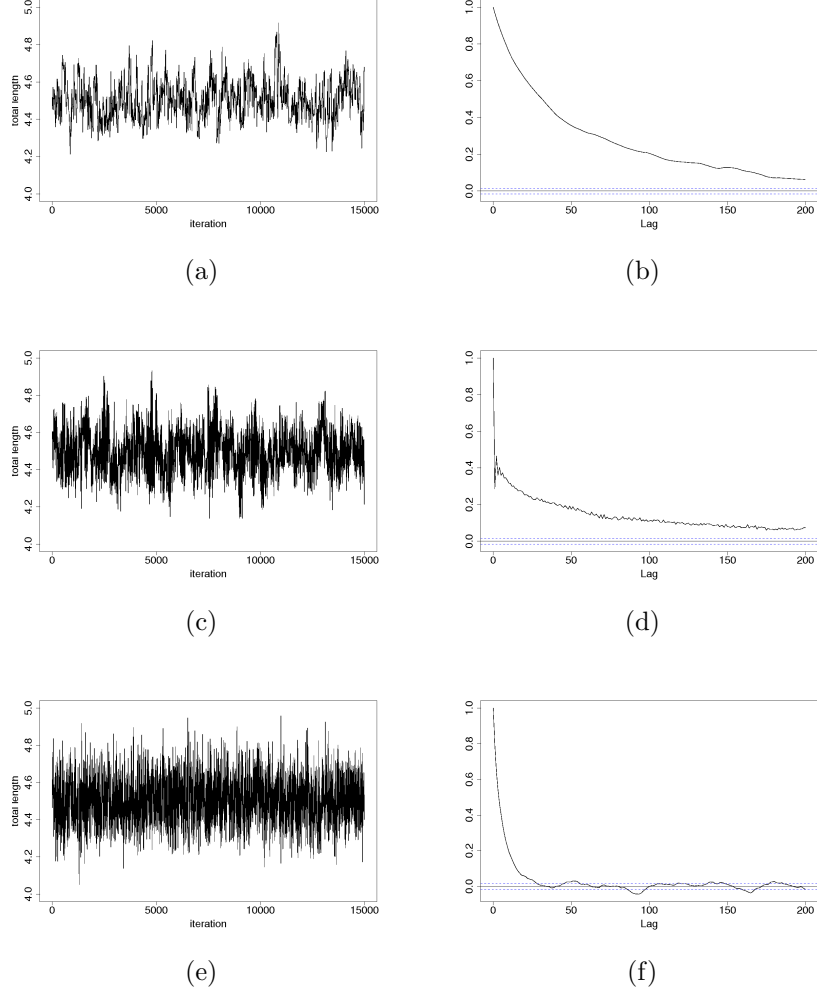


Figure 7: Trace plots of total branch length T , for 15 000 samples after burn-in. (a) MCMC with LOCAL proposal; (c) $(MC)^3$ with LOCAL proposal; (e) MCMC with NNI, BLNA, BLM updates. (b), (d), (f) Autocorrelation functions for samples of T output by (a), (c) and (e), respectively.

a few trees, the limited movement of the chain in topology space will restrict the movement in branch-length space, causing bad performance of the sampler.

Turning to the performance of the ϵ Dirichlet proposal for updating r and π , we considered the same synthetic alignment as before. To isolate the effects of the ϵ Dirichlet proposal, the topology and branch lengths were held fixed at their generating values and only π and r were updated. We monitored the r_{GT} chain.

Figures 8(a) and (b) show the traceplot and autocorrelation function, respectively, of the post-burn-in samples generated from a MrBayes run effectively implementing the version of the ϵ Dirichlet proposal with $\epsilon = 0$ (Larget and Simon, 1999). It is clear that there are some mixing problems and this is reflected in the estimated integrated autocorrelation $\hat{\tau} = 474$ (with the run taking 51 seconds).

As a computational comparison and a demonstration that this is not a rare event, figures 8(e) and (f) show the equivalent plots with the same proposal implemented in our own code; here $\hat{\tau} = 137$ (taking 354 seconds). It is possible to avoid the poor mixing by deploying (MC)³. Figures 8(c) and (d) show the traceplot and autocorrelation function, respectively, for MrBayes' coupled version. Coupling with one additional chain improves the mixing in this example ($\hat{\tau} = 46$) but at the increased computational cost from the original 51 seconds to 92 (although this doubling of time is clearly worthwhile for a reduction of 10 in $\hat{\tau}$).

Finally, Figures 8(g) and (h) are plots of r_{GT} updated with the ϵ Dirichlet proposal using $\epsilon = 1 \times 10^{-5}$. The stickiness at zero observed with the non-shifted Dirichlet proposal is avoided; here $\hat{\tau} = 52$ at virtually no extra computational cost, 366 seconds as opposed to the $\epsilon = 0$ cost of 354 seconds with our code. This demonstrates that our ϵ Dirichlet proposal is considerably cheaper in terms of computational cost relative to a tempered MCMC solution, and that it effectively removes the stickiness of the chain at values close to zero.

In order to assess the sensitivity of the ϵ -corrected proposal to the choice of ϵ , we performed a number of runs for a range of values for ϵ (keeping all other tuning parameters and initial values unchanged). Table 2 reports the ergodic averages of rate r_{GT} and the estimated integrated autocorrelation times, $\hat{\tau}$, corresponding to 15 000 samples after burn-in. The integrated autocorrelation times indicate that the worst-performing samplers are those with $\epsilon = 0$ (corresponding to the non-corrected proposal) and $\epsilon > 0.002$ (corresponding to a case where ϵ is greater than the generating r_{GT}). The poor performance of the latter is

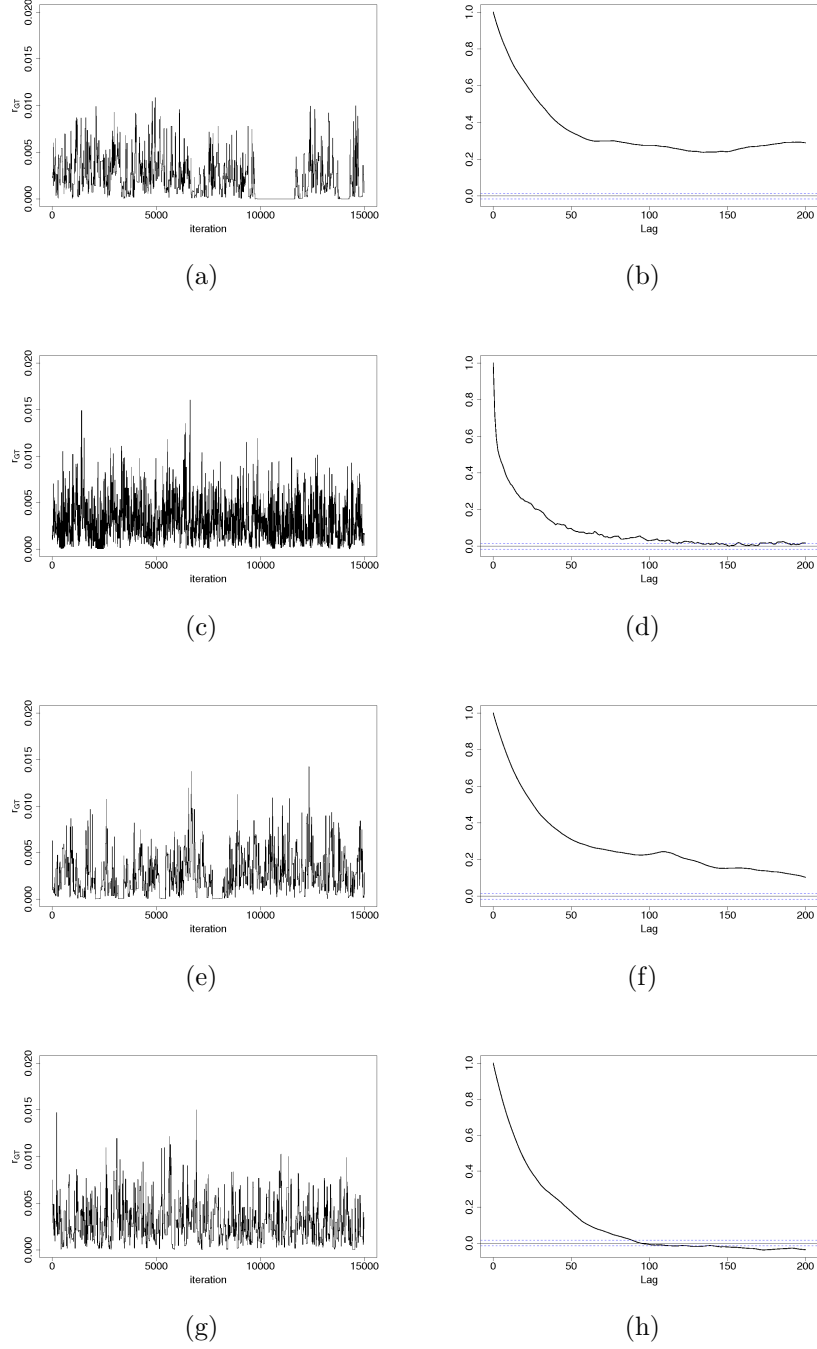


Figure 8: Traceplots of r_{GT} , for 15 000 samples after burn-in. (a) MrBayes with Dirichlet proposal; (c) MrBayes (MC)³ with Dirichlet proposal; (e) authors' code with Dirichlet proposal; and (g) authors' code with ϵ Dirichlet proposal. (b),(d),(f),(h) Autocorrelation functions of sampled r_{GT} , generated with methods (a), (c), (e) and (g), respectively. In all cases, tuning parameter α was set to 500.

ϵ	\bar{r}_{GT}	$\hat{\tau}(r_{GT})$
0	0.0027	137
1×10^{-8}	0.0028	71
1×10^{-7}	0.0029	70
1×10^{-6}	0.0029	80
1×10^{-5}	0.0029	52
1×10^{-4}	0.0029	60
1×10^{-3}	0.0028	68
2×10^{-3}	0.0029	66
3×10^{-3}	0.0027	156

Table 2: Ergodic averages for rate r_{GT} (generating value 0.002) and estimated integrated autocorrelation times $\hat{\tau}$ for 15 000 samples after burn-in.

because the proposal mechanism makes sufficiently large steps away from the bulk of the posterior support that many are rejected. Between these two extremes, there is considerable stability which dispels any potential worries over the estimation of very small rates. Overall, this shifted Dirichlet proposal seems a rather efficient approach.

6 Discussion

We have presented a classification method for molecular sequence data that employs mixture models augmented with allocation variables. Our method contrasts with a common approach that assigns different phylogenetic models to individual, *a priori*-known partitions of the data, later combining the partitions into a single composite model (e.g. Nylander et al., 2004). Partitioning the data prior to analysis makes two key assumptions: (1) the classes are known, and (2) there are as many classes as partitions in the alignment, which may result in overparametrisation. By contrast, our classification approach via mixture models enables us to discover the most appropriate segmentation of sites conditional on the model, and to make statements about the most important factors of substitutional heterogeneity in the data (e.g. is it the branch lengths that drive the segmentation in the data or the pattern

of substitution rates?).

Our classification method is directly relevant to the investigator who has a concatenation of multiple genes but not prior information about which of these genes can be appropriately described by the same parameters. It may equally be used as a tool for gene profiling: when having a set of reference genes with known functions and a query gene with an unknown profile, our method can be used to study the relation between the query gene and the reference genes. This should allow for statements about evolutionary similarities or dissimilarities between the reference and the query genes, leading to a clearer picture of the query gene's evolutionary identity.

Our method accounts for both qualitative and quantitative among-site rate variation by considering mixtures with multiple sets of branch lengths and Q matrices. Mixtures with multiple sets of branch lengths further account for a phenomenon known as heterotachy (Lopez, Casane and Philippe, 2002), in which the rates of evolution along branches leading to different taxa in the tree vary across sites. In fact, since the beginning of this research a number of groups have independently proposed mixtures of sets of branch lengths as a way of modelling heterotachy in phylogenetic studies (eg. Pagel and Meade, 2008).

An area of further work is the development of a mixture of overall-rates of substitution. An overall-rate for the j th component acts as a scaling factor of the branch lengths so that different components have trees with proportionally scaled branch lengths. Such a formulation would allow for a direct comparison with the popular discrete-gamma model (Yang, 1994), in which the overall-rates are assumed to conform to a Gamma distribution and all the components are constrained to have equal relative sizes. In contrast, a mixture of overall-rates allows more flexibility by letting the data support different relative sizes of the components and by avoiding distributional assumptions on the overall-rates; it further enables site classification by the inclusion of allocation variables, a feature that the discrete-gamma model lacks. However, the advantages of a mixture would have to be weighted against the convenient fact that a discrete-gamma formulation incorporates only one extra parameter relative to a homogeneous model. At present we are not aware of any study that has systematically compared the discrete-gamma model with a mixture on overall-rates and we believe that this is an interesting area of future research.

A further potential application of our classification method is as a tool for identifying

the sites that are unable to undergo substitution. The presence of invariant sites is a well-documented cause of inconsistency in phylogeny reconstruction (e.g. Steel et al., 2000), and site classification could be employed to pinpoint the invariant sites that should be excluded from the alignment before inference. This idea was previously discussed in Huelsenbeck and Suchard (2007), and it would require to define a mixture that includes a strictly invariant class (i.e. a class with total tree length constrained to being zero).

The MCMC methods discussed in this paper have been coded in a C program and source files are available upon request.

A Inference on allocation variables

Our MCMC sampler simulates a Markov chain of mixture parameters and a chain of allocation variables from the joint posterior distribution of all model parameters. Once the chains have been produced and checked for convergence to stationarity, good mixing and lack of label-switching, they can be used to make reliable inferences about the posterior distribution. In this appendix we describe ways of summarising the chain of allocation variables.

Consider a sample $z_n^{(1)}, \dots, z_n^{(M)}$ of the allocation for site n , generated from an MCMC run of length M after burn-in. Variable $z_n^{(i)}$ indicates the identity of the component to which site n is allocated at iteration i and it takes values in the set $\{1, \dots, k\}$. The sample $z_n^{(1)}, \dots, z_n^{(M)}$ can be used to count the number of times that site n was allocated to component j throughout the run. This frequency count, divided by the total number of samples, M , gives the *posterior classification probability* of site n to component j .

B Estimation of integrated autocorrelation times

The integrated autocorrelation time, $\tau(f)$, provides the means to compare different MCMC methods: to optimise the accuracy of estimation one could choose a method with the smallest possible $\tau(f)$. However, when comparing two methods, we must not only consider the accuracy of estimation but also the computational cost. In our implementation, we use Geyer (1992) initial positive sequence estimator to estimate $\tau(f)$:

$$\hat{\tau}(f) = -1 + 2 \sum_{i=0}^K \hat{\Gamma}_i,$$

where $\hat{\Gamma}_i = \hat{\rho}_{2i} + \hat{\rho}_{2i+1}$ is the sum of adjacent pairs of sample autocorrelations and $\hat{\rho}_t$ is the autocorrelation at lag t . Here K is chosen to be the largest integer such that $\hat{\Gamma}_i > 0$ for $i = 0, 1, \dots, K$.

The integrated autocorrelation time encodes the information about the correlation structure of the chain; the greater the correlation between the samples the larger the $\tau(f)$. It is in this sense that $\tau(f)$ is a measure of the ability of the chain to move agilely around the support of the posterior distribution. A rapidly moving (or mixing) chain produces more reliable estimates than a slowly mixing one (for a fixed number of iterations), and the former is usually preferred over the latter as long as the computational cost of the rapidly-mixing chain is not prohibitive.

Acknowledgements

Elisa Loza-Reyes is grateful to the Mexican Council for Science and Technology (CONACYT) for their financial support. We wish to thank Klaus Kurtenbach, Gabi Margos and Ed Feil, all from the Department of Biology and Biochemistry at the University of Bath, for their helpful discussions. Thanks also to Kevin Dawson, from Rothamsted Research, for valuable comments to earlier versions of this manuscript. This research was supported by the Bath Institute for Complex Systems (EPSRC grant GR/S86525/01).

References

- [1] W. M. Brown, E. M. Prager, A. Wang, and A. C. Wilson. Mitochondrial DNA sequences of primates: tempo and mode of evolution. *Journal of Molecular Evolution*, 18:225–239, 1982.
- [2] Joseph Felsenstein. Evolutionary trees from DNA sequences: A maximum likelihood approach. *Journal of Molecular Evolution*, 17:368–376, 1981.

- [3] Joseph Felsenstein and G. A. Churchill. A hidden Markov model approach to variation among sites in rate of evolution. *Molecular Biology and Evolution*, 13(1):93–104, 1996.
- [4] W. M. Fitch and E. Markowitz. An improved method for determining codon variability in a gene and its application to the rate of fixation of mutations in evolution. *Biochemical Genetics*, 4:579–593, 1970.
- [5] C. J. Geyer. Markov-chain Monte-Carlo maximum-likelihood. In Keramidas, E. M., editor, *Computing Science and Statistics*, pages 156–163, 1991.
- [6] C. J. Geyer. Practical Markov Chain Monte Carlo. *Statistical Science*, 7(4):473–511, 1992.
- [7] P. J. Green and X.-L. Han. Metropolis methods, Gaussian proposals, and antithetic variables. *Lecture Notes in Statistics*, 74:142–164, 1992.
- [8] M. Hasegawa, H. Kishino, and T. Yano. Dating of the human-ape splitting by a molecular clock of mitochondrial DNA. *Journal of Molecular Evolution*, 22:160–174, 1985.
- [9] K. Hayasaka, T. Gojobori, and S. Horai. Molecular phylogeny and evolution of primate mitochondrial DNA. *Molecular Biology and Evolution*, 5(6):626–644, 1988.
- [10] J. P. Huelsenbeck and F. Ronquist. MrBayes: Bayesian inference of phylogeny. *Bioinformatics*, 17:754–755, 2001.
- [11] J.P. Huelsenbeck and M.A. Suchard. A Nonparametric Method for Accommodating and Testing Across-Site Rate Variation. *Systematic Biology*, 56(6):975–987, 2007.
- [12] M. A. Hurn, P. J. Green, and F. Al-Awadhi. A Bayesian hierarchical model for photometric redshifts. *Applied Statistics, Journal of the Royal Statistical Society C*, 57:487–504, 2008.
- [13] T. Jukes and C. Cantor. Evolution of protein molecules. In H. N. Munro, editor, *Mammalian Protein Metabolism*, pages 21–132. Academic Press, New York, USA, 1969.
- [14] R. E. Kass and A. E. Raftery. Bayes Factors. *Journal of the American Statistical Association*, 90(430):773–795, 1995.

- [15] B. Kolaczowski and J. W. Thornton. A mixed branch length model of heterotachy improves phylogenetic accuracy. *Molecular Biology and Evolution*, 25(6):1054–1066, JUN 2008.
- [16] C. Lakner, P. van der Mark, J. P. Huelsenbeck, B. Larget, and F. Ronquist. Efficiency of Markov chain Monte Carlo tree proposals in Bayesian phylogenetics. *Systematic Biology*, 57(1):86–103, 2008.
- [17] C. Lanave, G. Preparata, C. Saccone, and G. Serio. A new method for calculating evolutionary substitution rates. *Journal of Molecular Evolution*, 20:86–93, 1984.
- [18] B. Larget and D. L. Simon. Markov chain Monte Carlo algorithms for the Bayesian analysis of phylogenetic trees. *Molecular Biology and Evolution*, 16(6):750–759, 1999.
- [19] N. Lartillot and H. Philippe. A Bayesian Mixture Model for Across-Site Heterogeneities in the Amino-Acid Replacement Process. *Molecular Biology and Evolution*, 21(6):1095–1109, 2004.
- [20] D. S. Leslie. Discussion on Model-based clustering for social networks (by M. S. Handcock, A. E. Raftery and J. M. Tantrum). *Journal of the Royal Statistical Society A*, 170:301–354, 2007.
- [21] P. Lopez, D. Casane, and H. Philippe. Heterotachy, an important process of protein evolution. *Molecular Biology and Evolution*, 19(1):1–7, JAN 2002.
- [22] A. Meade and M. Pagel. A Phylogenetic Mixture Model for Heterotachy. In P. Pontarotti, editor, *Evolutionary Biology from Concept to Application*, pages 29–41. Springer-Verlag, first edition, 2008.
- [23] G. W. Moore, M. Goodman, and J. Barnabas. An iterative approach from the standpoint of the additive hypothesis to the dendrogram problem posed by molecular data sets. *Journal of Theoretical Biology*, 38:423–457, 1973.
- [24] M.A. Newton and A.E. Raftery. Approximate Bayesian inference by the weighted likelihood bootstrap. *Journal of the Royal Statistical Society Series B*, 56:3–48, 1994.

- [25] J.A.A. Nylander, F. Ronquist, J.P. Huelsenbeck, and J.L. Nieves-Aldrey. Bayesian Phylogenetic Analysis of Combined Data. *Systematic Biology*, 53(1):47–67, 2004.
- [26] M. Pagel and A. Meade. A phylogenetic mixture model for detecting pattern-heterogeneity in gene sequence or character-state data. *Systematic Biology*, 56(4):571–581, 2004.
- [27] M. Pagel and A. Meade. Modelling heterotachy in phylogenetic inference by reversible-jump Markov chain Monte Carlo. *Philosophical Transactions of the Royal Society B*, 363:3955–3964, 2008.
- [28] A. Rambaut and N. C. Grassly. Seq-Gen: An application for the Monte Carlo simulation of DNA sequence evolution along phylogenetic trees. *Computer Applications in the Biosciences*, 13:235–238, 1997.
- [29] D. F. Robinson. Comparison of labeled trees with valency three. *Journal of Combinatorial Theory*, 11:105–119, 1971.
- [30] M. Steel, D. Huson, and P. J. Lockhart. Invariable Sites Models and Their Use in Phylogeny Reconstruction. *Systematic Biology*, 49:225–232, 2000.
- [31] M. A. Suchard, R. E. Weiss, and J. S. Sinsheimer. Bayesian selection of continuous-time Markov chain evolutionary models. *Molecular Biology and Evolution*, 18(6):1001–1013, 2001.
- [32] S. Tavaré. Some probabilistic and statistical problems in the analysis of DNA sequences. In R. M. Miura, editor, *Lectures on Mathematics in the Life Sciences*, volume 17, pages 57–86. American Mathematical Society, Providence, USA, 1986.
- [33] A. Webb, J. M. Hancock, and C. C. Holmes. Phylogenetic inference under recombination using Bayesian stochastic topology selection. *Bioinformatics*, 25(2):197–203, 2009.
- [34] W. Xie, P. O. Lewis, Y. Fan, and M-H. Chen. Improving Marginal Likelihood Estimation for Bayesian Phylogenetic Model Selection. *Systematic Biology*, 60:150–160, 2011.
- [35] Z. Yang. Maximum-likelihood estimation of phylogeny from DNA sequences when substitution rates differ over sites. *Molecular Biology and Evolution*, 10(6):1396–1401, 1993.

- [36] Z. Yang. Maximum likelihood phylogenetic estimation from DNA sequences with variable rates over sites: approximate methods. *Journal of Molecular Evolution*, 39:306–314, 1994.
- [37] Z. Yang. A space-time process model for the evolution of DNA sequences. *Genetics*, 139:993–1005, 1995.
- [38] Z. Yang. *Computational Molecular Evolution*. Oxford Series in Ecology and Evolution. Oxford University Press, Great Britain, 2006.
- [39] Z. Yang. PAML 4: Phylogenetic Analysis by Maximum Likelihood. *Molecular Biology and Evolution*, 24:1586–1591, 2007.

Supplementary Material: Alternating between BLNA and BLM, or using only one?

One of the advantages of an alternating BLNA&BLM mechanism for proposing branch lengths is that one move can be tuned to generate modest steps while the other to produce bolder ones. In this section, we investigate the performance of single BLNA or BLM updates, compared to a sampler that uses both of them in an alternated manner. To do so, we produced a synthetic DNA alignment of size 6 sequences \times 2500 sites, generated with the software package Seq-Gen (Rambaut and Grassly, 1997) under the GTR model of nucleotide substitution. The values used to generate this alignment are the following: the phylogenetic tree and branch lengths in Newick format¹ are: $((Taxon2 : 0.16, Taxon4 : 0.34) : 0.61, Taxon5 : 0.2) : 0.53, (Taxon3 : 0.48, Taxon6 : 2.14) : 0.35, Taxon1 : 0.05)$, where $\{Taxon1, \dots, Taxon6\}$ is the set of leaf-labels and the numbers correspond to the lengths of the branches; the vector of stationary probabilities is $\pi = (\pi_A = \dots = \pi_G = 0.25)$ and the substitution rates are $r = (r_{AC} = 0.140, r_{AG} = 0.340, r_{AT} = 0.090, r_{CG} = 0.008, r_{CT} = 0.420, r_{GT} = 0.002)$.

In this exercise, we fixed r , π and the tree to their true values. The target distribution is the joint posterior for branch lengths. We generated candidate branch lengths according to three different methods: (A) from a BLNA proposal; (B) from a BLM proposal; and (C) from an alternating BLNA&BLM scheme. In the alternating BLNA&BLM scheme, candidate branch lengths were generated from the BLNA proposal at even iterations and from the BLM proposal at odd ones. The justification for alternating moves and still converging to the target distribution is given by the fact that if chains P and R have the same stationary distribution, so does PR (see, for example, Grimmett and Stirzaker, 2004).

We produced three replicates under each method, varying the tuning parameters of the BLM move (δ parameter) and BLNA (σ parameter). The settings for these replicates are shown in Table 1.

Table 2 reports the ergodic averages and estimated integrated autocorrelation times for each of the six exterior branch lengths t_1, \dots, t_6 , based on 15 000 samples after a burn-in period of 5 000 iterations. We only report exterior branch lengths since interior branches are not uniquely labelled across different tree topologies. On the top-right of Table 2, we have indicated the replicate

¹The Newick format is widely used in phylogenetics for representing trees in computer-readable form. It makes use of the correspondence between trees and nested parentheses. This format is further described in Felsenstein, 2004.

<i>replicate</i>	δ	σ
(i)	1.1	0.08
(ii)	1.3	0.06
(iii)	1.5	0.04

Table 1: Values for the δ tuning parameter of the BLM move and the σ tuning parameter of the BLNA proposal. These parameters were varied for three different replicates, (i), (ii) and (iii).

number and the same order applies for all branches. Note the better performance of BLNA (small $\hat{\tau}$) relative to BLM in estimating $\mathbb{E}(t_1|x)$. The results suggest unsuitability of BLM for estimating short branches, which can be further investigated by calculating the expected value and variance of a branch length with respect to the BLM proposal.

In a BLM move, a candidate length is generated as $b' = bm$, where b is the current length and m is a random variable with density function $f(m) = (\lambda m)^{-1}$, $1/\delta < m < \delta$. The expected value and variance of b' are given by

$$\begin{aligned}
\mathbb{E}_f(b') &= b \mathbb{E}_f(m) \\
&= \frac{b}{\lambda} \left(\delta - \frac{1}{\delta} \right) \\
\text{Var}_f(b') &= b^2 \text{Var}_f(m) \\
&= b^2 \left(\frac{1}{2\lambda} \left(\delta^2 - \frac{1}{\delta^2} \right) - \frac{1}{\lambda^2} \left(\delta - \frac{1}{\delta} \right)^2 \right), \quad \frac{b}{\delta} < b' < \delta b
\end{aligned} \tag{1}$$

where $\lambda = 2 \log(\delta)$ and $\delta > 1$ is a tuning parameter. In the limit $b \rightarrow 0$, the expected value and the variance of the candidate length approach $\mathbb{E}_f(b') \rightarrow 0$ and $\text{Var}_f(b') \rightarrow 0$. This produces a phenomenon in which the chain is unable to move away from the zero neighbourhood, which we have dubbed ‘zero-stickiness’. A phenomenon like this results in poor estimation performance, since the chain spends several iterations trapped at a small neighbourhood of the state space, producing MCMC samples that are highly correlated to one another.

On the other hand, a candidate branch length is generated from a BLNA proposal as $b' = b + \sigma u$, where $u \sim N(0, 1)$ and $\sigma > 0$ is the tuning parameter. Under this proposal, the variance of b' does not depend on the current branch length and the step-size of the move is not influenced but by σ .

	<i>true length</i>	(A) <i>BLNA</i>		(B) <i>BLM</i>		(C) <i>BLNA&BLM</i>	
		<i>average</i>	$\hat{\tau}$	<i>average</i>	$\hat{\tau}$	<i>average</i>	$\hat{\tau}$
t_1	0.05	0.070	11.01	0.071	38.01	0.069	14.06 (i)
		0.069	14.18	0.066	251.84	0.070	18.30 (ii)
		0.070	13.36	0.070	24.41	0.071	13.63 (iii)
t_2	0.16	0.159	10.88	0.157	9.80	0.157	11.76
		0.158	17.28	0.157	26.73	0.156	19.65
		0.158	9.69	0.157	9.73	0.157	8.83
t_3	0.48	0.469	17.40	0.468	22.53	0.469	25.53
		0.467	21.41	0.465	56.76	0.470	27.70
		0.470	29.28	0.470	26.95	0.469	30.80
t_4	0.34	0.310	10.05	0.313	10.05	0.312	11.01
		0.311	16.52	0.313	19.88	0.312	14.39
		0.311	11.53	0.312	11.28	0.312	10.52
t_5	0.20	0.188	9.50	0.188	11.61	0.187	11.08
		0.188	10.65	0.188	35.64	0.188	17.04
		0.187	10.53	0.189	9.87	0.187	12.20
t_6	2.14	2.074	36.17	2.076	17.21	2.069	21.56
		2.081	25.97	2.071	18.12	2.072	20.39
		2.066	76.32	2.074	19.89	2.075	21.73

Table 2: The ergodic averages for exterior branch lengths and the estimated integrated autocorrelation time ($\hat{\tau}$), for three different branch-length updating methods: (A) BLNA proposal; (B) BLM proposal; and (C) an alternating BLNA&BLM scheme. For each method, three replicates were performed, each replicate with different tuning parameters. The replicate number is indicated on the top right-hand-side of the table, and this same order applies for all branches. The results correspond to 15 000 samples after burn-in. All runs were initialised at the same starting point. The average execution time (across replicates) for the three methods were: (A) 3 300, (B) 3 320, and (C) 3 000; all measured in seconds.

In estimating $\mathbb{E}(t_6|x)$, BLNA performs poorly relative to BLM (see Table 2). In other words, when the true branch length is long, BLM outperforms BLNA in all replicates. We believe that this might be due to the fact that the step-size of BLM depends on the current branch length whereas that of BLNA does not. A method that alternates between BLNA and BLM inherits the good properties of both proposals. The results for the combined BLNA&BLM, in Table 2, highlight the good estimation performance of such a strategy and justify our preference for alternating between BLNA and BLM when updating branch lengths.

References

- [1] Joseph Felsenstein. *Inferring Phylogenies*. Sinauer Associates Inc., USA, 2004.
- [2] G. R. Grimmett and D. R. Stirzaker. *Probability and Random Processes*. Oxford University Press, Great Britain, third edition, 2004.
- [3] A. Rambaut and N. C. Grassly. Seq-Gen: An application for the Monte Carlo simulation of DNA sequence evolution along phylogenetic trees. *Computer Applications in the Biosciences*, 13:235–238, 1997.